

LIBRARY
OF THE
MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

.M414
no. 211-
66

A Multivariate Adaptive Control Model

211-66

John D. C. Little

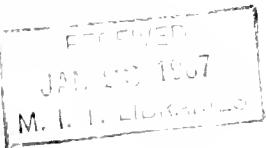
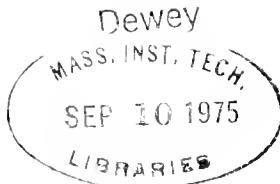


A Multivariate Adaptive
Control Model

211-66

John D. C. Little

August 18, 1966



Abstract

A specialized multivariate adaptive control model is developed. Each of r control variables is to be set in each of a sequence of time periods. The process being controlled has a response (profit) function that is the sum of a constant plus linear and quadratic forms in the control variables. The coefficients of the quadratic form are assumed to be known constants, those of the linear form to change with time as first order, autoregressive processes. Information about the changing coefficients is collected by performing a 2^r factorial experiment on a subportion of the process being controlled. Provision is also made for adding further information from unspecified sources. Bayesian methods are used to update distributions of the unknown coefficients. The values of the control variables are set to maximize the sum of discounted future profits, as are the experimental design parameters. The probabilistic assumptions of the model are chosen so that all distributions are normal with known variances and, for the most part, zero covariances between variables. Partly as a result of this, optimal control turns out to involve rather simple exponential smoothing rules.

1. Introduction

Reference 1 presents an adaptive control model that has the following characteristics.

- (1) A process is to be controlled by setting a single control variable in each of a sequence of time periods.
- (2) The profit or response of the process is a quadratic function of the control variable.
- (3) The coefficient of the linear term of the response function is imperfectly known and changes with time as a first order, autoregressive process. The coefficient of the quadratic term is a known constant.
- (4) An experiment is performed in each time period to gain information about the linear coefficient.
- (5) The optimal setting of the control variable is shown to be given by an exponential smoothing process.
- (6) The optimal experimental design parameters are relatively easy to determine. They are chosen to minimize the sum of two expected costs: the cost of operating with poor response information and the cost of deliberate non-optimal operation used in the experiment.
- (7) The distributions of all random variables involved are normal with known variances.

The present paper generalizes the work in the following ways:

- (1) The number of control variables is increased from 1 to r.
- (2) The profit function becomes the sum of linear and quadratic

forms in the r variables. The coefficients of the linear form change with time as a first order, autoregressive process; those of the quadratic form are known constants.

- (3) The experiment becomes a 2^r factorial design.
- (4) Provision is made for the introduction of information about the process beyond that developed in the experiment.

As before, the optimal setting of the control variables and the optimal experimental design parameters are found. Dynamic programming arguments are used to justify the optimization results in somewhat more detail than in [1].

2. Profit Model

Let $\underline{x} = [x_1, \dots, x_r]^T$ = a vector of control variables
 $p(\underline{x})$ = profit

Profit is a random variable. When we wish to emphasize that a quantity is a random variable, we shall place a tilde (\sim) over it. We shall assume

$$\begin{aligned}\tilde{p}(\underline{x}) &= \tilde{\alpha} + \tilde{\beta}^T \underline{x} - \underline{x}^T \underline{\Gamma} \underline{x} \\ &= \tilde{\alpha} + \sum_{j=1}^r \tilde{\beta}_j x_j - \sum_{j=1}^r \sum_{k=1}^r x_j \gamma_{jk} x_k\end{aligned}\tag{2.1}$$

where

$$\begin{aligned}\tilde{\beta} &= [\tilde{\beta}_1, \dots, \tilde{\beta}_r]^T = \text{a vector of imperfectly known coefficients,} \\ \tilde{\alpha} &= \text{an imperfectly known constant,} \\ \underline{\Gamma} &= [\gamma_{ij}] = \text{a matrix of perfectly known constants}\end{aligned}$$

We presume that the conditions of the application dictate that (2.1) has a maximum for some reasonable \underline{x} . A mathematical assumption that guarantees the existence of a maximum for $\underline{x} \in \mathbb{R}^r$ is that $\underline{\Gamma}$ be positive

definite. We shall assume this, and further, without loss of generality, that $\underline{\Gamma}$ is symmetric. Under these circumstances $\underline{\Gamma}^{-1}$ exists and is also positive definite symmetric. Let

$$\theta_{ij} = ij \text{ element of } \underline{\Gamma}^{-1}.$$

Necessary conditions for maximizing $p(\underline{x})$ over $\underline{x} \in R^r$ are

$$\frac{\partial p}{\partial x_j} = \beta_j - 2 \sum_{k=1} \gamma_{jk} x_k = 0 \quad j = 1, \dots, r.$$

The unique solution to these equations is

$$\underline{\tilde{x}}^* = (1/2) \underline{\Gamma}^{-1} \underline{\beta}. \quad (2.2)$$

(2.2) maximizes p as may be seen from the following. Let

$$\begin{aligned} \ell(\underline{x}) &= p(\underline{x}^*) - p(\underline{x}) \\ &= \underline{\beta}^T (\underline{x}^* - \underline{x}) - \underline{x}^{*T} \underline{\Gamma} \underline{x}^* + \underline{x}^T \underline{\Gamma} \underline{x} \\ &= (\underline{x}^* - \underline{x})^T \underline{\Gamma} (\underline{x}^* - \underline{x}) \end{aligned} \quad (2.3)$$

$\underline{\Gamma}$ positive definite implies $y^T \underline{\Gamma} y > 0$ for all $y \neq 0$. Thus, $\ell(\underline{x})$ is minimized with value zero when $\underline{x} = \underline{x}^*$ and, by (2.3), $p(\underline{x})$ is maximized.

Notice that

$$\begin{aligned} \tilde{\ell}(\underline{x}) &= (\tilde{\underline{x}}^* - \underline{x})^T \underline{\Gamma} (\tilde{\underline{x}}^* - \underline{x}) \\ &= \text{loss relative to perfect information.} \end{aligned} \quad (2.4)$$

All the variables and constants so far, with exception of $\underline{\Gamma}$ will change with time. Time periods will be indexed by t . When we wish to make the time dependence explicit we shall write $\underline{x}(t)$, $\underline{\beta}(t)$, etc., but oftentimes we shall omit t to reduce clutter.

3. Information Changes

The principal hindrance to making good decisions is imperfect knowledge of $\underline{\beta}(t)$. We shall first preview the information changes that take place for $\underline{\beta}(t)$ and establish notation for them.

Let $\tilde{z} \sim f(\cdot)$

stand for "the random variable \tilde{z} has the probability density function $f(\cdot)$." Let

$f_N(\cdot | \underline{\mu}, \underline{h})$ = normal density function with vector of means $\underline{\mu}$
and precision matrix \underline{h} .

$$\tilde{\underline{\beta}} = (\tilde{\alpha}, \tilde{\beta}_1, \dots, \tilde{\beta}_r).$$

Notice that we have added $\tilde{\alpha}$ to the $\tilde{\underline{\beta}}$ vector. This will be done in the present section only. The value of α does not directly affect optimal decision making since it is the constant of the profit function. However, α is involved in the statistics because information about α usually implies information about the β_j and vice versa.

Let

$\underline{b}^{(i)} = [a^{(i)}, b_1^{(i)}, \dots, b_r^{(i)}]^T$ = a vector of statistics about $\tilde{\underline{\beta}}$
at the i^{th} step of information
change within a time period.

$\underline{h}^{(i)} = [\underline{h}_{ij}^{(i)}]$ = a matrix of precisions at the
 i^{th} step

$\underline{\gamma}^{(i)} = [\gamma_{ij}^{(i)}] = [\underline{h}^{(i)}]^{-1}$ = covariance matrix corresponding
to $\underline{h}^{(i)}$.

Like $\underline{\beta}$, $\underline{b}^{(i)}$, $\underline{h}^{(i)}$, and $\underline{\gamma}^{(i)}$ will have an extra dimension corresponding to α in this section but not elsewhere.

The cycle of information stages for a time period is:

(1) Start. At the beginning of t , just before making the decision on $\underline{x}(t)$, the knowledge of $\underline{\beta}(t)$ will be represented by

$$\tilde{\underline{\beta}}(t) \cap f_N(\cdot | \underline{b}^{(1)}(t), \underline{h}^{(1)}(t))$$

(2) Experiment. During t an experiment will be run that retrospectively will provide information about $\underline{\beta}(t)$. Sufficient statistics for the experimental results will be:

$$\underline{b}^{(2)} \text{ and } \underline{h}^{(2)}.$$

(3) Update. After the experimental results are absorbed, the (posterior) distribution of $\underline{\beta}(t)$ will be:

$$\tilde{\underline{\beta}}(t) \cap f_N(\cdot | \underline{b}^{(3)}, \underline{h}^{(3)})$$

(4) Change. A change, partially unpredictable, takes place in $\underline{\beta}$ between t and $t + 1$. This will be governed by a first order, autoregressive process mechanism with parameterization

$$\underline{k}, \underline{b}^{(4)}, \text{ and } \underline{h}^{(4)}$$

(5) Update. The change leaves us with information about $\underline{\beta}(t+1)$ summarized in

$$\tilde{\underline{\beta}}(t+1) \cap f_N(\cdot | \underline{b}^{(5)}, \underline{h}^{(5)})$$

(6) Further Information. Various activities not individually modeled may tell us in advance something about $\underline{\beta}(t+1)$. This will be summarized in sufficient statistics

$$\underline{b}^{(6)} \text{ and } \underline{h}^{(6)}$$

(7) Update. Assimilation of the outside information gives

$$\tilde{\underline{b}}(t+1) \cap f_N(\cdot | \underline{b}^{(7)}, \underline{h}^{(7)}).$$

Identifying

$$\underline{b}^{(1)}(t+1) = \underline{b}^{(7)}$$

$$\underline{h}^{(1)}(t+1) = \underline{h}^{(7)}$$

we are ready to recycle.

The steps are next traced in detail.

3.1 Start. At the beginning of t , certain prior information about $\underline{b}(t)$ is available. This will be the information used to decide the value of $\underline{x}(t)$. We shall assume that $\underline{b}(t)$ has a multivariate normal density with mean vector $\underline{b}^{(1)}(t)$ and known precision matrix $\underline{h}^{(1)}(t)$ and that $\underline{h}^{(1)}(t)$ is diagonal.

$$\tilde{\underline{b}}(t) \cap f_N(\cdot | \underline{b}^{(1)}(t), \underline{h}^{(1)}(t))$$

$$E[\tilde{\underline{b}}] = \underline{b}^{(1)} = [a^{(1)}, b_1^{(1)}, \dots, b_r^{(1)}]$$

$$V[\tilde{\underline{b}}] = \underline{v}^{(1)} = [\underline{h}^{(1)}]^{-1}$$

$$\underline{h}^{(1)} = \begin{bmatrix} \frac{1}{M} & & & \\ \hline & 0 & & \\ 0 & h_{11}^{(1)} & & 0 \\ & 0 & & h_{rr}^{(1)} \end{bmatrix}$$

$$M = \text{large number.}$$

In assuming M large, we are assuming that the prior information on α is negligible. Since the precision matrix is diagonal, the covariance matrix is too, with non-zero elements simply the reciprocals of the corresponding precisions.

3.2 Experiment. Since β is imperfectly known and subject to unpredictable changes, its value will be monitored over time. The method of doing this will be a 2^r factorial experiment conducted each period. The experiment is presumed to be conducted on a portion of the actual process being controlled. Thus, if the process occupies an area (as in agriculture or marketing) a number of subareas might be split off for experimental treatment. For flow processes (as in chemical operations), the basic time period might be subdivided so that, for certain intervals, given experimental treatments are run. We shall assume that

N = total number of experimental units into which the process can be divided

n = number of experimental units devoted to the 2^r factorial experiment

The experiment will be designed in terms of deviations from the "normal" operations conducted in the $N-n$ non-experimental units. Each control variable will have levels designated "high" and "low". If normal operations call for setting the j^{th} control variable to x_j , the experimental levels will be:

$$\text{high: } x_j + \Delta_j/2 \quad \text{and}$$

$$\text{low: } x_j - \Delta_j/2.$$

In a complete factorial design, each possible combination of experimental levels is applied to some experimental unit.

Let $\underline{\delta}_i \equiv [\delta_{i1}, \dots, \delta_{ir}]^T$ = the vector of deviations for experimental unit i (thus, δ_{ij} equals either $+\Delta_j/2$ or $-\Delta_j/2$ as appropriate for i).

$\underline{x}_i = \underline{x} + \underline{\delta}_i$ = the vector of control variables applied to experimental unit i.

The statistical model for the profit observation from the i^{th} experimental unit is

$$\tilde{p}(\underline{x}_i) = [1, \underline{x}_i^T] \underline{\beta} - \underline{x}_i^T \underline{\Gamma} \underline{x}_i + \tilde{\epsilon}_i, \quad i = 1, \dots, n,$$

where $\underline{\beta}$ is fixed for the given t but imperfectly known and

$$\tilde{\epsilon}_i \sim f_N(\cdot | 0, 1/\sigma^2).$$

The variance, σ^2 , for the experimental units will be assumed known. Since $\underline{\Gamma}$ and \underline{x} are known, we can simplify notation by defining

$$\tilde{y}_i = \tilde{p}(\underline{x}_i) + \underline{x}_i^T \underline{\Gamma} \underline{x}_i$$

Then data point i can be represented by the linear model

$$\tilde{y}_i = [1, \underline{x}_i^T] \underline{\beta} + \tilde{\epsilon}_i \quad (3.2.1)$$

Equation (3.2.1) may be regarded as a multiple regression with α and β_i as the unknown coefficients. The Bayesian analysis of a normal regression with known variance is given by Raiffa and Schlaifer [3] (p. 334 ff). Let

$$\underline{y} = [y_1, \dots, y_n]^T$$

$$\underline{\underline{X}} = \begin{bmatrix} 1 & \underline{\underline{x}}_1^T \\ \vdots & \vdots \\ \vdots & \vdots^T \\ 1 & \underline{\underline{x}}_n^T \end{bmatrix}$$

Sufficient statistics for the experiment are:

$$\underline{\underline{h}}^{(2)} \equiv (1/\sigma^2) \underline{\underline{X}}^T \underline{\underline{X}} = (\underline{\underline{y}}^{(2)})^{-1}$$

$$\underline{\underline{b}}^{(2)} \equiv (1/\sigma^2) \underline{\underline{y}}^{(2)} \underline{\underline{X}}^T \underline{\underline{y}}$$

Because of the special form of the $\underline{\underline{x}}_i$ in the factorial design,

$$\underline{\underline{h}}^{(2)} = (n/\sigma^2) \left[\begin{array}{c|cccc} 1 & \underline{\underline{x}}_1 & \underline{\underline{x}}_2 & \dots & \underline{\underline{x}}_r \\ \hline \underline{\underline{x}}_1 & \underline{\underline{x}}_1^2 + \Delta_1^2/4 & \underline{\underline{x}}_1 \underline{\underline{x}}_2 & & \underline{\underline{x}}_1 \underline{\underline{x}}_r \\ \underline{\underline{x}}_2 & \underline{\underline{x}}_2 \underline{\underline{x}}_1 & \underline{\underline{x}}_2^2 + \Delta_2^2/4 & & \underline{\underline{x}}_2 \underline{\underline{x}}_r \\ \vdots & & \ddots & & \vdots \\ \vdots & \underline{\underline{x}}_r \underline{\underline{x}}_1 & \underline{\underline{x}}_r \underline{\underline{x}}_2 & \underline{\underline{x}}_r^2 + \Delta_r^2/4 & \end{array} \right] \quad (3.2.2)$$

Whence

$$\underline{\underline{y}}^{(2)} = (\sigma^2/n) \left[\begin{array}{c|cccc} 1 + \sum_{j=1}^r 4\underline{\underline{x}}_j^2 / \Delta_j^2 & -4\underline{\underline{x}}_1 / \Delta_1^2 & -4\underline{\underline{x}}_2 / \Delta_2^2 & \dots & -4\underline{\underline{x}}_r / \Delta_r^2 \\ \hline -4\underline{\underline{x}}_1 / \Delta_1^2 & 4/\Delta_1^2 & 0 & & 0 \\ -4\underline{\underline{x}}_2 / \Delta_2^2 & 0 & 4/\Delta_2^2 & & 0 \\ -4\underline{\underline{x}}_r / \Delta_r^2 & 0 & 0 & \ddots & 4/\Delta_r^2 \end{array} \right]$$

If the components of $\underline{\underline{b}}^{(2)}$ are denoted $\underline{\underline{b}}^{(2)} = [a^{(2)}, b_1^{(2)}, \dots, b_r^{(2)}]$,

we find that

$$a^{(2)} = (1/n) \sum_{i=1}^n y_i - \sum_{j=1}^r b_j^{(2)} \underline{\underline{x}}_j$$

$$b_j^{(2)} = (2/n\Delta_j) \{ \sum_{i \in I_j^+} y_i - \sum_{i \in I_j^-} y_i \}$$

where

$$I_j^+ = \{i \mid \delta_{ij} = +\Delta_j/2\}$$

$$I_j^- = \{i \mid \delta_{ij} = -\Delta_j/2\}$$

The components of $\underline{b}^{(2)}$ are the classical estimates of the regression coefficients for this data. Notice in $\underline{\gamma}^{(2)}$ that the covariance matrix for the $b_j^{(2)}$ is diagonal so that information we collect about an individual β_j in the experiment is independent of the other β_j 's and, subject to possible operational restrictions, is of independently controllable accuracy through Δ_j .

3.3 Update. The experimental information on $\tilde{\beta}(t)$ is summarized in $\underline{b}^{(2)}$ and $\underline{h}^{(2)}$ and will be combined with prior information in the manner shown in [3]. Denoting the posterior distribution by

$$\tilde{\beta}(t) \sim f_N(\cdot \mid \underline{b}^{(3)}, \underline{h}^{(3)}),$$

we have

$$\underline{h}^{(3)} = \underline{h}^{(2)} + \underline{h}^{(1)} = [\underline{\gamma}^{(3)}]^{-1},$$

$$\underline{b}^{(3)} = \underline{\gamma}^{(3)} [\underline{h}^{(1)} \underline{b}^{(1)} + \underline{h}^{(2)} \underline{b}^{(2)}].$$

Adding (3.1.1) and (3.2.2) and neglecting the term in $1/M$ compared to 1, we get $\underline{h}^{(3)}$. Taking its inverse gives:

$$\underline{\underline{v}}^{(3)} = \begin{vmatrix} (\sigma^2/n) + \sum_{j=1}^r v_{jj}^{(3)} x_j^2 & -v_{11}^{(3)} x_1 & -v_{22}^{(3)} x_2 & \cdots & -v_{rr}^{(3)} x_r \\ -v_{11}^{(3)} x_1 & v_{11}^{(3)} & 0 & 0 & 0 \\ -v_{22}^{(3)} x_2 & 0 & v_{22}^{(3)} & \ddots & 0 \\ \vdots & & & \ddots & \ddots \\ -v_{rr}^{(3)} x_r & 0 & 0 & \ddots & v_{rr}^{(3)} \end{vmatrix}$$

where the diagonal terms are given by:

$$v_{jj}^{(3)} = [(n\Delta^2)/(4\sigma^2) + h_{jj}^{(1)}]^{-1} \quad j = 1, \dots, r.$$

In other words, the posterior variance of β_j is the reciprocal of the sum of the prior precision and what might be called the experimental precision. Of central importance to later mathematical simplicity is the result that $v_{jj}^{(3)}$ is independent of $\underline{x}(t)$.

Working out $\underline{b}^{(3)} = [a^{(3)}, b_1^{(3)}, \dots, b_r^{(3)}]^T$, we obtain

$$a^{(3)} = a^{(2)} + \sum_{j=1}^r x_j(b_j^{(2)} - b_j^{(3)})$$

$$b_j^{(3)} = v_{jj}^{(3)} h_{jj}^{(1)} b_j^{(1)} + v_{jj}^{(3)} [(n\Delta^2)/(4\sigma^2)] b_j^{(2)} \quad j = 1, \dots, r$$

Essentially $b_j^{(3)}$ is formed by weighting $b_j^{(1)}$ and $b_j^{(2)}$ proportionally to their respective precisions.

3.4 Change. The process being controlled is changing with time.

Some of the changes may be improvements instituted by the controlling organization. Some may be the result of process elements deteriorating

slowly or abruptly. We model them all as a random change in β between t and $t + 1$. The result is a deterioration of information, although we shall later give the controller an opportunity to perceive the new β , if perhaps imperfectly.

The model of change will be a first order, autoregressive process.

Given that $\tilde{\beta}(t) = \beta(t)$:

$$\tilde{\beta}(t+1) = k \underline{\beta}(t) + (\underline{I} - k \underline{b})^{(4)} + \tilde{\epsilon}_{\beta}(t+1) \quad (3.4.1)$$

where

$$k = \begin{bmatrix} 0 & 0 & \cdots & 0 \\ \hline 0 & k_1 & \cdots & 0 \\ & \ddots & \ddots & \\ 0 & 0 & \cdots & k_r \end{bmatrix}$$

$$\tilde{\epsilon}_{\beta}(t) \sim f_N(\cdot | 0, \underline{h}^{(4)})$$

$$[\underline{h}^{(4)}]^{-1} = \underline{v}^{(4)} = \begin{bmatrix} M_1 & 0 & \cdots & 0 \\ \hline 0 & \sigma_1^2 & \cdots & 0 \\ & \ddots & \ddots & \\ 0 & 0 & \cdots & \sigma_r^2 \end{bmatrix}, M_1 \text{ large,}$$

$$\underline{b}^{(4)} = [a^{(4)}, b_1^{(4)}, \dots, b_r^{(4)}]^T$$

Thus, the individual β_i change independently. The value of a β_i at $t+1$ is a weighted combination of the value at t and a long run average value $b_i^{(4)}$ plus a random change having zero mean and variance σ_i^2 . The a term is treated specially. It is picked newly in each time period from a distribution with mean, $a^{(4)}$, and a large variance, M_1 . In effect, the information on a is lost each time period.

3.5 Update. The next step is to find the distribution of $\tilde{\beta}(t+1)$.

Since it is a linear combination of multivariate normal random variables, $\tilde{\beta}(t+1)$ is itself multivariate normal. Therefore, it suffices to find its mean vector, $\underline{b}^{(5)}$ and its covariance matrix, $\underline{\Sigma}^{(5)}$ in order to specify its distribution.

From (3.4.1), $\tilde{\beta}(t+1)$ has mean and variance:

$$\underline{b}^{(5)} = \underline{k}^{(3)} \underline{b}^{(4)} + (\underline{I} - \underline{k}) \underline{b}^{(4)}$$

$$= [a^{(4)}, k_1 b_1^{(3)} + (1-k_1) b_1^{(4)}, \dots, k_r b_r^{(3)} + (1-k_r) b_r^{(4)}]^T$$

$$\underline{\Sigma}^{(5)} = \underline{k}^{(3)} \underline{k}^T + \underline{\Sigma}^{(4)}$$

$$= \begin{bmatrix} M_1 & 0 & \dots & 0 \\ 0 & k_1^2 \underline{\Sigma}_{11}^{(3)} + \sigma_1^2 & & \\ \vdots & & \ddots & \\ \vdots & & & \\ 0 & 0 & & k_r^2 \underline{\Sigma}_{rr}^{(3)} + \sigma_r^2 \end{bmatrix}$$

3.6 Further Information. We may suppose that the controller of the process learns about it in more ways than the experiment, particularly since the controller himself may be making some of the changes. We shall model the new information as a multinormal process of the same form used earlier to treat the experimental information. The perceived value of $\tilde{\beta}(t+1)$, given that in fact $\tilde{\beta}(t+1) = \beta(t+1)$, will be denoted $\underline{b}^{(6)}$. Its likelihood function will be assumed normal:

$$\underline{b}^{(6)} \cap f_N(\cdot | \underline{\beta}(t+1), \underline{h}^{(6)})$$

We shall take

$$\underline{h}^{(6)} = \left[\begin{array}{c|ccc} 1/M_2 & 0 & & 0 \\ \hline 0 & \ell_1 & \dots & 0 \\ & 0 & \ddots & \\ 0 & & & \ddots & \ell_r \end{array} \right], \quad M_2 \text{ large},$$

Thus we assume that no new information is gained about α and that the new information does not interconnect the β_i .

3.7 Update. The updating is done as in the case of the experiment.

The posterior for $\underline{\beta}(t+1)$ has precision matrix and mean vector:

$$\begin{aligned} \underline{h}^{(7)} &= \underline{h}^{(6)} + \underline{h}^{(5)} \\ &= \left[\begin{array}{c|ccccc} 1/k_1 + 1/M_2 & 0 & \dots & & 0 \\ \hline 0 & \ell_1 + h_{11}^{(5)} & & & 0 \\ \vdots & & \ddots & & \\ 0 & 0 & & \ddots & \ell_r + h_{rr}^{(5)} \end{array} \right] \\ \underline{b}^{(7)} &= \underline{y}^{(7)} [\underline{h}^{(6)} \underline{b}^{(6)} + \underline{h}^{(5)} \underline{b}^{(5)}]. \end{aligned}$$

The components of $\underline{b}^{(7)}$ are

$$\begin{aligned} a^{(7)} &= [M_1 a^{(6)} + M_2 a^{(5)}] / (M_1 + M_2) \\ b_j^{(7)} &= [h_{jj}^{(6)} / h_{jj}^{(7)}] b_j^{(6)} + [h_{jj}^{(5)} / h_{jj}^{(7)}] b_j^{(5)} \quad j = 1, \dots, r \end{aligned}$$

We are now ready to recycle. We set $\underline{b}^{(1)}(t+1) = \underline{b}^{(7)}$ and $\underline{h}^{(1)}(t+1) = \underline{h}^{(7)}$. Since $\tilde{\beta}(t+1)$ is multivariate normal with a diagonal precision matrix and

a small precision for α , the starting condition assumptions have been preserved.

It may be well to recapitulate and summarize in preparation for the next section. We have moved from a prior on $\underline{\beta}(t)$ characterized by $\underline{b}^{(1)}(t)$ and diagonal $\underline{\underline{h}}^{(1)}(t)$ to a prior on $\widetilde{\beta}(t+1)$ characterized by $\underline{b}^{(1)}(t+1)$ and diagonal $\underline{\underline{h}}^{(1)}(t+1)$. The diagonal terms of $\underline{\underline{h}}^{(1)}(t+1)$ may be written in terms of $\underline{\underline{h}}^{(1)}(t)$ as follows:

$$\begin{aligned} h_{00}^{(1)}(t+1) &= (1/M_1) + (1/M_2) \\ h_{jj}^{(1)}(t+1) &= \ell_j + \frac{1}{\sigma_j^2 + \frac{k_j^2}{(n\Delta_j^2/4\sigma_j^2) + h_{jj}^{(1)}(t)}} \quad j = 1, \dots, r \end{aligned} \quad (3.7.1)$$

All the constants in this recursion are known in advance: ℓ_j is a precision in the "further information" process; σ_j^2 is a variance from the change process; k_j^2 is a smoothing constant from the change process; and $(4\sigma_j^2/n\Delta_j^2)$ is an experimental variance.

The recursion for $\underline{b}^{(1)}(t+1)$ in terms of $\underline{b}^{(1)}(t)$ is:

$$\begin{aligned} a^{(1)}(t+1) &= [M_1 a^{(6)} + M_2 a^{(5)}] / (M_1 + M_2) \\ b_j^{(1)}(t+1) &= [h_{jj}^{(6)} / h_{jj}^{(1)}(t+1)] b_j^{(6)} + [h_{jj}^{(5)} / h_{jj}^{(1)}(t+1)] \{ \\ &\quad (1-k_j) b_j^{(4)} + k_j \{ [\nu_{jj}^{(3)} / \nu_{jj}^{(2)}] b_j^{(2)} + [\nu_{jj}^{(3)} h_{jj}^{(1)}(t)] b_j^{(1)}(t) \} \} \\ &\quad j = 1, \dots, r \end{aligned} \quad (3.7.2)$$

Thus, $b_j^{(1)}(t+1)$ is a linear combination of various $b_j^{(i)}$. Of these $b_j^{(4)}$ is a fixed constant, the long term average β_j in the change process.

The others are developed during the cycle of information changes: $b_j^{(2)}$
 is the experimental observation and $b_j^{(6)}$ is the "further information"
 data. The weighting factors are various combinations of variances
 and precisions, all known constants.

Of critical importance in the next section is the fact that neither $\underline{h}^{(1)}(t+1)$ nor $\underline{b}^{(1)}(t+1)$ depend on $\underline{x}(t)$.

4. Optimization

There are two decision problems: how to set the control variables and how to determine the experimental design parameters. We shall proceed by dynamic programming to solve these problems. The objective function will be the discounted value of the expected loss compared to perfect information. As mentioned earlier, we drop the α components in $\underline{\beta}$, $\underline{b}^{(1)}$, and $\underline{h}^{(1)}$ now that we are out of Section 3.

4.1 Formulation. As shown earlier the loss relative to perfect

information in any time period, t , is:

$$\tilde{\ell}(\underline{x}) = (\underline{x} - \tilde{\underline{x}}^*)^T \Gamma (\underline{x} - \tilde{\underline{x}}^*)$$

where

$$\tilde{x}^* = (1/2)\Gamma^{-1} \tilde{\beta}(t) \quad (4.1.1)$$

Each time period is entered with information about $\underline{b}(t)$ summarized in $\underline{b}^{(1)}(t)$ and $\underline{h}^{(1)}(t)$. These statistics will be the state variables of the dynamic program. The superscript (1) will be dropped for brevity.

Also, following dynamic programming convention, time will be indexed backward. To take note of the change, we index with s instead of t .

Let

$f_s(\underline{b}(s), \underline{h}(s))$ = minimum expected loss over the last s periods.

$f_0(\underline{b}(0), \underline{h}(0)) = 0$

The loss over the last $s+1$ periods will be made up of (1) losses in period $s+1$ and (2) the loss in the remaining s periods. The losses in $s+1$ consist of a loss attributable to imperfect choice of \underline{x} , an operating loss from using part of the process to experiment on, and possibly an out-of-pocket cost of experimentation. Let

c = out-of-pocket cost per experimental unit (e.g. special data collection costs)

ρ = discount factor

The dynamic programming recursion is:

$$\begin{aligned} f_{s+1}(\underline{b}(s+1), \underline{h}(s+1)) &= \min_{n, \Delta_1, \dots, \Delta_r} \min_{\underline{x}} E[(N-n)(\underline{x} - \underline{\tilde{x}}^*)^T \Gamma (\underline{x} - \underline{\tilde{x}}^*) \\ &\quad + \sum_{i=1}^n (\underline{x} + \underline{\delta}_i - \underline{\tilde{x}}^*)^T \Gamma (\underline{x} + \underline{\delta}_i - \underline{\tilde{x}}^*) + nc \\ &\quad + \rho f_s(\underline{b}(s), \underline{h}(s))] \end{aligned} \quad (4.1.2)$$

The $\underline{b}(s)$ and $\underline{h}(s)$ are related to $\underline{b}(s+1)$ and $\underline{h}(s+1)$ and to the experimental constants by (3.7.1) and (3.7.2) after change of indexing.

Several simplifications can be made. Because of the symmetries of the experimental design

$$\begin{aligned} \sum_{i=1}^n (\underline{x} + \delta_i - \underline{x}^*)^T \underline{\Gamma} (\underline{x} + \delta_i - \underline{x}^*) \\ = n(\underline{x} - \underline{x}^*)^T \underline{\Gamma} (\underline{x} - \underline{x}^*) + (n/4) \sum_{j=1}^r \gamma_{jj} \Delta_j^2 . \end{aligned}$$

Furthermore, making use of (4.1.1), we can obtain

$$\begin{aligned} E[(\underline{x} - \tilde{x}^*)^T \underline{\Gamma} (\underline{x} - \tilde{x}^*)] &= (\underline{x} - E(\tilde{x}^*))^T \underline{\Gamma} (\underline{x} - E(\tilde{x}^*)) \\ &\quad + E[(\tilde{x}^* - E(\tilde{x}^*))^T \underline{\Gamma} (\tilde{x}^* - E(\tilde{x}^*))] \\ &= (\underline{x} - E(\tilde{x}^*))^T \underline{\Gamma} (\underline{x} - E(\tilde{x}^*)) + (1/4) \sum_{j=1}^r \theta_{jj} / h_{jj} \end{aligned}$$

Both $\gamma_{jj} > 0$ and $\theta_{jj} > 0$ for all j . This follows from $\underline{\Gamma}$ positive definite, for if any $\gamma_{jj} \leq 0$, we could choose $x_j > 0$ and $x_k = 0$, $k \neq j$, and contradict the definition of positive definite. Since $\underline{\Gamma}^{-1}$ is positive definite too, the same remark holds for its diagonal elements θ_{jj} .

The recursion (4.1.2) now becomes

$$\begin{aligned} f_{s+1}(\underline{b}(s+1), \underline{h}(s+1)) &= \min_{n, \Delta_1, \dots, \Delta_r} \min_{\underline{x}} \{ N(\underline{x} - E(\tilde{x}^*))^T \underline{\Gamma} (\underline{x} - E(\tilde{x}^*)) \\ &\quad + (n/4) \sum_{j=1}^r \gamma_{jj} \Delta_j^2 + (N/4) \sum_{j=1}^r \theta_{jj} / h_{jj}(s+1) + nc \\ &\quad + \rho E[f_s(\underline{b}(s), \underline{h}(s))] \} \end{aligned} \quad (4.1.3)$$

4.2 Optimal Setting of Control Variables. As noted earlier, neither $\underline{b}(s)$ nor $\underline{h}(s)$ depends on \underline{x} , nor does any term on the right of (4.1.3) except the first. Since $\underline{\Gamma}$ is positive definite, the optimal $\underline{x}(s+1)$ is

$$\underline{x}(s+1) = E(\tilde{\underline{x}}^*) = (1/2) \underline{\Gamma}^{-1} E(\tilde{\underline{b}}(s+1)) = (1/2) \underline{\Gamma}^{-1} \underline{b}(s+1)$$

or, in terms of earlier notation:

$$\underline{x}(t) = (1/2) \underline{\Gamma}^{-1} \underline{b}^{(1)}(t). \quad (4.2.1)$$

This solves the problem of how to set the control variables at t .

Substitution of this \underline{x} into (4.1.3) removes the first term and the min over \underline{x} . Observe next that, since $f_0 = 0$, f_1 is not a function of $\underline{b}(1)$, and, by induction, f_s is not a function of $\underline{b}(s)$. Therefore, we drop $\underline{b}(s)$ from the state description. Furthermore, since f_s contains no random variables the expectation operator can be dropped. The recursion becomes

$$f_{s+1}(\underline{h}(s+1)) = \min_{n, \Delta_1, \dots, \Delta_r} \left\{ (1/4) \sum_{j=1}^r \gamma_{jj} n \Delta_j^2 + (N/4) \sum_{j=1}^r \theta_{jj} / h_{jj}(s+1) \right. \\ \left. + nc + \rho f_s(\underline{h}(s)) \right\} \quad (4.2.2)$$

where $\underline{h}(s)$ is related to $\underline{h}(s+1)$ by (3.7.1), or, in the present notation

$$h_{jj}(s) = \ell_j + \frac{1}{\sigma_j^2 + \frac{k_j^2}{\frac{n \Delta_j^2}{4\sigma} + h_{jj}(s+1)}} \quad (4.2.3)$$

4.3 Optimal Experimental Design. In finding the experimental design parameters, two cases may be distinguished: $c = 0$ and $c > 0$, according as the cost of the experimental units is negligible or not.

We shall confine most of our attention to the $c = 0$ case and make a few remarks about the effect of $c > 0$ at the end. In addition, we shall restrict our attention to the situation where the parameters of the problem make it optimal to experiment on each variable in each time period, except possibly for the last few. This will not always be the case, since, for example, "further information" could be so good that no experimentation is necessary.

When $c = 0$, n and Δ_j appear only in the combination $n\Delta_j^2$, as may be seen in (4.2.2) and (4.2.3). Therefore, the quantities $n\Delta_j^2$ may be regarded as the decision variables. The exact values of n and Δ_j can then be partly determined by other conditions, e.g., n must be some multiple of 2^r to have a complete factorial design.

Next we observe that, since $f_0 = 0$, $f_1(h(l))$ will be a sum of terms each involving variables and parameters in j alone. The minimization separates into a minimization for each j . By induction the same will be true for any s . Therefore we can write

$$f_s(h(s)) = \sum_{j=1}^r g_{sj}(h_{jj}(s)) \quad (4.3.1)$$

where

$$g_{s+1,j}(h_{jj}(s+1)) = \min_{n\Delta_j^2} \{ (\gamma_{jj}/4) n\Delta_j^2 + N\theta_{jj}/4h_{jj}(s+1) + \rho g_{sj}(h_{jj}(s)) \} \quad (4.3.2)$$

$$g_{0j} = 0 ,$$

and (4.2.3) gives $h_{jj}(s)$ in terms of $h_{jj}(s+1)$ and $n\Delta_j^2$. All other quantities are known constants.

The dynamic programming problem (4.3.2) has been solved by Hurst [2] in the case $\ell_j = 0$. Our work will be based on his but we shall here deduce only some of his results. First define a new decision variable

$$z_j(s+1) = (\eta \Delta_j^2) / (4\sigma^2) + h_{jj}(s+1) \quad (4.3.3)$$

Operationally, z_j is the precision of $\tilde{\beta}_j$ after the experiment has been performed. Since the starting precision $h_{jj}(s+1)$ is known, z_j determines $\eta \Delta_j^2$. If an experiment is to be performed on variable j , $z_j(s+1) > h_{jj}(s+1)$.

Next (4.3.2) is rewritten in terms of z_j and, in addition, two time periods of the recursion are displayed. Thus

$$\begin{aligned} g_{s+1,j}(h_{jj}(s+1)) = \min_{z_j(s+1) \geq h_{jj}(s+1)} & \{ \gamma_{jj} \sigma^2 [z_j(s+1) - h_{jj}(s+1)] \\ & + N\theta_{jj}/4h_{jj}(s+1) + \rho \min_{z_j(s) \geq h_{jj}(s)} \{ \gamma_{jj} \sigma^2 [z_j(s) - h_{jj}(s)] \\ & + N\theta_{jj}/4h_{jj}(s) + \rho f_{s-1}(h_{jj}(s-1)) \} \} \end{aligned} \quad (4.3.4)$$

$$\text{where } h_{jj}(s) = \ell_j + z_j(s+1)/(\sigma_j^2 z_j(s+1) + k_j^2) \quad (4.3.5)$$

All the terms in (4.3.4) that depend on $z_j(s+1)$ have been boxed, provided that $z_j(s)$ is not affected by the constraint $z_j(s) \geq h_{jj}(s)$, i.e. provided that an experiment is performed at s under optimal operation. This is just the case we are considering. Observe that none of the boxed terms are affected by $z_j(s)$ so that the minimization over $z_j(s)$ will not

be affected by our choice of $z_j(s+1)$, i.e. future decisions will not be affected by the present one. Under these circumstances we can split off the $z_j(s+1)$ decision. We proceed by solving

$$\min_{z_j(s+1) \geq 0} \{ \gamma_{jj}^{-\sigma} [z_j(s+1) - \rho h_{jj}(s)] + \rho N \theta_{jj} / 4 h_{jj}(s) \} \quad (4.3.6)$$

where $h_{jj}(s)$ is given by (4.3.5). This is a relatively simple one-dimensional minimization problem involving the ratio of polynomials. Notice that $z_j(s+1)$ is determined in terms of the constants of the process and is good for any s (except for the last few, where the recursion (4.3.2) becomes somewhat special). Therefore we shall denote the solution of (4.3.6) simply as z_j . Substitution of z_j into (4.3.5) gives the next period's starting $h_{jj}(s)$ which again will be the same for all large s and so will be denoted h_{jj} .

To recapitulate, the optimal experiment parameter $n_{\Delta j}^2$ is found as follows. Suppose we are at time $s+1$ with precision $h_{jj}(s+1)$. We solve (4.3.6) for z_j , and set $n_{\Delta j}^2$ for $s+1$ to

$$n_{\Delta j}^2 = 4\sigma^2 [z_j - h_{jj}(s+1)].$$

Use of this z_j ensures that in s we will have the h_{jj} found by putting z_j into (4.3.5). Thereafter, for all large s , $n_{\Delta j}^2$ will be constant at

$$n_{\Delta j}^2 = 4\sigma^2 [z_j - h_{jj}].$$

If $z_j \leq h_{jj}(s+1)$ or $z_j \leq h_{jj}$, we have uncovered a case where it is not optimal to experiment in every time period.

This takes care of choosing the design parameters when $c = 0$ and it is optimal to experiment. When $c > 0$ and it is optimal to experiment, the smallest feasible $n > 0$ is best with the given cost structure. This is because Δ_j can always be set to achieve the optimal $n\Delta_j^2$. In some situations, however, there may be practical constraints on n or Δ_j .

5. Discussion

Certain assumptions bear further comment. We have assumed that the quadratic constants of the profit function, $\underline{\Gamma}$, are known. This seems to be a very strong assumption. Notice, however, that the 2^r factorial experiment actually gathers information about $\underline{\Gamma}$ through the second order interaction terms. Thus, even if $\underline{\Gamma}$ were poorly known at the start, one would expect to learn it over a period of time.

It has been assumed that α has high period to period variance. Although α does not directly enter the setting of the control variables, good knowledge of it would provide a way to learn β better. Our assumption dismisses this possibility and thereby gains for us considerable mathematical simplification. In some processes α probably does have high variance. The virtue of cross section analyses, for example, is to eliminate this source of variability. In classical experimental design it is often desirable to run different treatments in the same time period because absolute response (i.e. α) is subject to changes with time that may quite obscure relative response (i.e. β).

In summary, a rather complex, r-variable adaptive control model has been solved. Optimal control involves setting control variables by linear combinations of observed quantities and known constants. The measurement process in the control procedure involves an experiment whose parameters are determined largely through simple one dimensional minimizations.

To a considerable degree, the simplicity of the optimal solution is the result of careful selection of assumptions. Nevertheless, we suggest that the idealized model offers insight into the nature of good solutions for situations where the assumptions do not exactly hold. Sensitivity analyses of the one-variable model in [1] indicate that this may often be the case.



References

1. J. D. C. Little, "A Model of Adaptive Control of Promotional Spending," to appear in Operations Research, December 1966.
2. E. G. Hurst, Jr., Memorandum, July 17, 1966.
3. H. Raiffa and R. Schlaifer, Applied Statistical Decision Theory, Graduate School of Business Administration, Harvard University, Boston, 1961.



Date Due

DEC 19 '75	FEB 19 '76	
DEC 22 '75		
JAN 20 '76		
FEB 18 '76		
DEC 27 '76		
SPR 15 '77		
MAY 21 '78		
JUN 10 '78		
MAY 25 '79		
MAY 21 '82		



MIT LIBRARIES

3 9080 003 869 895

175000.00
HD28 M.I.T. Alfred P. Sloan
.M414 School of Management
Nos.211-66 Working Papers.
Nos.224-66 Nos.211-66 to Nos.224-66



SIGN this card and present it with book
at the CIRCULATION DESK.

MASSACHUSETTS INSTITUTE OF
TECHNOLOGY LIBRARIES

224-66

